IN THE CLAIMS:

1. (Currently Amended) In a file server having a storage operating system, a method for managing storage of data in a plurality of storage devices, each storage device having a plurality of blocks for storing data, comprising:

generating block layout information in a file system layer of the storage operating system by determining which blocks within the plurality of blocks are allocated for storing data and which are unallocated;

transferring the block layout information from the file system layer to a RAID layer of the storage operating system; ~~and~~

responsive to the block layout information, controlling ~~the~~ execution of I/O operations at the RAID layer by identifying a plurality of contiguous blocks on a single storage device within the plurality of blocks for use by each I/O operation so as to substantially maximize chain lengths of read operations for calculation of parity;

determining whether a parity subtraction method or a recalculation method requires a fewest number of read operations to calculate parity for the I/O operations;

selecting the~~a~~ parity subtraction method or the~~a~~ recalculation method for parity calculation based on ~~the~~ which method ~~that~~ requires the~~a~~ fewest number of read operations ~~to compute parity for the I/O operations~~; and

responsive to the block layout information and the parity calculation method selected, identifying the contiguous blocks within the plurality of blocks for use by the I/O operations.

2. (Currently Amended) A method for managing storage of data in a plurality of storage devices, each comprising a plurality of storage blocks, comprising:

generating block layout information;

determining whether a first methodology or a second methodology requires a fewest number of read operations to calculate parity; and

in response to the block layout information and the determination, controlling execution of I/O operations by identifying a plurality of contiguous storage blocks on a single storage device for use by each I/O operation so as to substantially minimize ~~a~~ the

number of read operations needed for calculation of error correction parameters across a stripe disposed among the plurality of storage devices.

3. (Original) The method of claim 2 wherein the calculation of error correction parameters comprises the calculation of parity.

4. (Currently Amended) The method of claim ~~3~~ 2 wherein the calculation of parity comprises selecting a parity calculation operation from a group consisting of a subtraction method as the first methodology and a parity re-calculation method as the second methodology.

5. (Currently Amended) The method of claim 2 wherein the identification of contiguous storage blocks for use in the I/O operation substantially maximizes a chain length by substantially maximizing ~~the~~ a number of blocks having a contiguous physical layout on the single storage device.

6. (Currently Amended) The method of claim 2, further comprising:

identifying the contiguous storage blocks for use in the I/O operation so as to substantially maximize the chain length by substantially maximizing ~~the~~ a number of blocks having sequential volume block numbers (VBNs) associated with the plurality of storage blocks.

7. (Currently Amended) The method of claim 2, further comprising:

identifying the contiguous storage blocks for use in the I/O operation so as to substantially maximize the chain length by substantially maximizing locality of the contiguous blocks of the single storage device.

8. (Currently Amended) The method of claim 2 wherein ~~the~~ controlling execution ~~step~~ comprises:

examining storage blocks to which data is to be written prior to write operations; and

selecting ~~one of a plurality of parity calculation methodologies including~~ either, ~~a~~ the first methodology comprising minimizing a number of blocks read, ~~and~~ or ~~a~~ the second methodology comprising maximizing chain lengths of blocks read for calculating parity ~~calculation~~.

9. (Currently Amended) The method of claim 8, wherein ~~the~~ controlling execution ~~step~~ further comprises:

implementing selection of the parity calculation methodology responsive to the block layout information; ~~and~~

wherein, if the selection constitutes ~~substantially~~ minimizing the number of blocks read,

determining on a stripe-by-stripe basis whether to calculate parity based on a subtraction method or a recalculation method,

performing any appropriate read operations to support the method selected, and

calculating parity responsive to the number of blocks read and the data to be written; and

wherein, if the selection constitutes ~~substantially~~ maximizing chain lengths of blocks read,

deciding which storage blocks to read to ~~substantially~~ maximize chain length while ~~substantially~~ minimizing the number of ~~storage~~ blocks read to support either ~~a~~ the subtraction method or ~~a~~ the recalculation method,

performing read operations on the number of blocks ~~to be~~ read, and

calculating parity responsive to the number of blocks read and the data to be written.

10. (Currently Amended) The method of claim 2, wherein ~~the identification of~~ identifying the contiguous storage blocks is based at least in part on an available resource.

11. (Original) The method of claim 2 further comprising transmitting the block layout information from a file system layer to a RAID layer.

12. (Currently Amended) The method of claim 2 wherein ~~the~~ generating ~~step~~ further comprises:

making a first determination as to whether a storage block is unallocated;

making a second determination as to a current implementation of the plurality of storage devices; and

generating the block layout information based at least in part on the first and the second determinations.

13. (Currently Amended) The method of claim 2, wherein the I/O operation is one of a plurality of I/O operations and at least one of the plurality of I/O operations is a read operation.

14. (Currently Amended) The method of claim ~~2~~ 5, wherein the chain length is a chain length of a read operation for calculation of parity.

15. (Currently Amended) The method of claim ~~2~~ 5, wherein the chain length is a chain length for a write operation for the data.

16. (Currently Amended) A method for managing storage of data in a storage system, comprising:

maintaining a plurality of storage devices each having a plurality of storage blocks; and

writing data to predetermined storage blocks of the plurality of storage blocks across a plurality of stripes and to predetermined contiguous storage blocks within each storage device so as to ~~substantially~~ maximize chain lengths of the predetermined contiguous storage blocks within each storage device and minimizing a number of read operations for the calculation of error correction parameters across each stripe of the

plurality of stripes by determining whether a parity subtraction method or a recalculation method requires a fewest number of read operations to calculate parity, and selecting thea parity subtraction method or thea recalculation method for parity calculation based on the which method that requires the fewest number of read operations to compute parity.

17. - 38. (Cancelled)

39. (Currently Amended) A storage system, comprising:

a plurality of storage devices each having a plurality of storage blocks; and

a storage manager in communication with the plurality of storage devices, the storage manager configured to write writing data to predetermined storage blocks across a plurality of stripes and to predetermined storage blocks within each storage device so as to substantially maximize chain length of the plurality of storage blocks by selecting as many as contiguous storage blocks within a single storage device while substantially minimizing a number of read operations required for calculation of error correction parameters across each stripe of the plurality of stripes by determining whether a parity subtraction method or a recalculation method requires a fewest number of read operations to calculate parity and selecting athe parity subtraction method or athe recalculation method for parity calculation based on the which method that requires the fewest number of read operations to compute parity.

40. (Currently Amended) A system for managing the storage of data, the system comprising:

a plurality of storage devices each having a plurality of storage blocks;

a storage device manager in communication with the plurality of storage blocks;

a block layout information generator in communication with the storage device manager and the plurality of storage blocks; and

an error correction parameter calculator in communication with the plurality of storage blocks and the storage device manager,

wherein the storage device manager, in response to the block layout information from the block layout information generator, controls ~~the~~ execution of an I/O operation by identifying a plurality of contiguous storage blocks on a single storage device for use by the I/O operation so as to ~~substantially~~ maximize chain length within the single storage device while ~~substantially~~ minimizing ~~the~~ a number of read operations required for calculation by the error correction parameter calculator of error correction parameters across a stripe by determining whether a parity subtraction method or a recalculation method requires a fewest number of read operations to calculate parity for the I/O operations and selecting ~~a~~ the parity subtraction method or ~~a~~the recalculation method for parity calculation based on ~~the~~ which method ~~that~~ requires the fewest number of read operations ~~to compute parity~~.

41 - 44. (Cancelled)

45. (Previously Presented) A method for managing storage of data by a server, comprising:

receiving a request to write the data to a plurality of storage devices;

generating block layout information to determine which blocks within a plurality of blocks located in the plurality of storage devices are allocated ~~for storing data~~ and which are unallocated;

identifying blocks within the plurality of blocks for use by a set of I/O operations ~~to store the data~~;

determining ~~the~~ a number of read operations needed to ~~compute~~ calculate parity for the data by ~~computing~~ calculating parity using a subtraction method ~~of computing parity~~;

determining ~~the~~ a number of read operations needed to ~~compute~~ calculate parity for the data by ~~computing~~ calculating parity using a recalculation method ~~of computing parity~~;

choosing either the subtraction method of ~~computing~~ calculating parity or the recalculation method of ~~computing~~ calculating parity by determining which ~~of these two~~

method~~s~~ requires ~~the~~ a fewer number of read operations, and choosing the which method that requires ~~requiring~~ the fewer number of read operations; and

writing the data to the identified blocks, and ~~computing~~ calculating parity for the data using the ~~chosen~~ which method ~~of computing parity~~.

46. (Previously Presented) The method of claim 45, further comprising:

choosing to either ~~firstly~~ maximize chain lengths of read operations for ~~calculation of~~ calculating parity or ~~secondly~~ choosing to place the data with a high degree of locality in the plurality of storage devices, by choosing the which method ~~which~~ that requires the ~~fewest~~ fewer number of read operations ~~in computing parity for the data~~.

47. (Currently Amended) A method for managing storage of data by a server, comprising:

receiving a request to write data to a plurality of storage devices;

generating block layout information to determine which blocks within a plurality of blocks located in the plurality of storage devices are allocated ~~for storing data~~ and which are unallocated;

identifying blocks within the plurality of blocks for use by a set of I/O operations to store the data; and

~~selecting~~ determining, in response to the block layout information, whether to ~~substantially~~ minimize ~~the~~ a number of read blocks or whether to ~~substantially~~ maximize chain lengths of read blocks based on which method requires a fewer number of read operations, and

implementing ~~the~~ a selection, during the writing of the data to the plurality of storage devices, responsive to the block layout information, and responsive to whether ~~substantially~~ minimizing the number of read blocks or substantially maximizing chain lengths of read blocks requires the fewer number of read operations.

48. (Currently Amended) The method of claim 47, further comprising:

in response to selecting to ~~substantially~~ minimize the number of read blocks, determining whether to calculate parity based on ~~the~~ a subtraction method or ~~the~~ a recalculation method ~~by determining~~ based on which method requires the fewer number of read operations~~, and selecting the method which requires the fewer number of read operations.~~; and

performing the write operation and calculating the parity using ~~the~~ which parity calculation method ~~requiring~~ requires the fewer number of read operations.

49. (Currently Amended) The method of claim ~~48~~ 47, further comprising:

in response to selecting to ~~substantially~~ maximize chain lengths of read blocks, deciding which storage blocks to read to ~~substantially~~ maximize chain lengths while minimizing the number of storage blocks read to support either the subtraction method or the recalculation method ~~of parity calculation~~; and

performing the write operation and calculating the parity using ~~the~~ which parity calculation method ~~requiring~~ requires the fewer number of read operations.

50. (Currently Amended) A method for managing storage of data by a server, comprising:

receiving a request to write data to a plurality of storage devices;

generating block layout information to determine which blocks within a plurality of blocks located in the plurality of storage devices are allocated ~~for storing data~~ and which are unallocated;

identifying blocks within the plurality of blocks for use by a set of I/O operations to store the data;

testing to either maximize chain lengths of read operations for calculation of parity, or to place the data with a high degree of locality in the plurality of storage devices, the testing ~~having the steps~~ comprising,

determining, for both maximizing chain length and placing the data with ~~a~~ the high degree of locality, ~~the~~ a number of read operations needed to ~~compute~~ calculate

parity for the data, by ~~computing~~ calculating parity using both ~~the~~ a subtraction method of ~~computing~~ calculating parity and ~~the~~ a recalculation method of ~~computing~~ calculating parity;

~~firstly~~ first choosing to either maximize chain lengths of read operations for calculation of parity or to place the data with ~~a~~ the high degree of locality in the plurality of storage devices, and after ~~this~~ the first choice, secondly choosing either the subtraction method of ~~computing~~ calculating parity or the recalculation method of ~~computing~~ calculating parity by determining which of these methods requires ~~the~~ a fewest number of read operations,

choosing ~~the~~ which method ~~requiring~~ requires the fewest number of read operations of ~~computing~~ calculating parity of the data; and

writing the data to the identified blocks, and ~~computing~~ calculating parity for the data using the ~~chosen~~ which method ~~of computing parity~~.

51. (Currently Amended) A computer readable media, comprising:

said computer readable media containing instructions for execution on a processor for a method of managing storage of data in a plurality of storage devices, each storage device having a plurality of blocks for storing data, the method ~~having~~ comprising,

generating block layout information; and

in response to the block layout information, controlling ~~the~~ execution of an I/O operation by identifying a plurality of contiguous storage blocks on a single storage device for use by the I/O operation so as to ~~substantially~~ minimize ~~the~~ a number of read operations needed for calculation of error correction parameters across a stripe by determining whether a parity subtraction method or a recalculation method requires a fewest number of read operations to calculate parity for the I/O operations and selecting ~~a~~the parity subtraction method or ~~a~~the recalculation method for parity calculation based on ~~the~~ which method ~~that~~ requires the fewest number of read operations ~~to compute parity~~.